DeepClean: A Deep Learning Model for Gravitational Wave Noise Reduction

Alan Mayer, Christian O'Connor, Ansh Thukral, and Haris Zia
University of California San Diego
(Dated: December 11, 2025)

Gravitational-wave detectors such as LIGO and Virgo have enabled direct observations of compact binary mergers, yet their sensitivity remains limited by environmental and instrumental noise. While classical noise-subtraction techniques perform well for linearly coupled disturbances, they struggle in the presence of nonlinear and nonstationary noise sources. Recent work has demonstrated that deep-learning–based approaches can overcome these limitations by learning complex couplings directly from auxiliary sensor data.

In this project, we investigate the DeepClean framework, a convolutional neural network designed to subtract environmental noise from gravitational-wave strain data using witness channels. Due to limited public availability of raw auxiliary data, we evaluate DeepClean using a simulation-based approach in which realistic noise sources—including broadband noise, power mains interference, and jitter noise—are algorithmically coupled into cleaned strain data. We assess model performance on both linearly and nonlinearly coupled noise scenarios, exploring the effects of different loss-function weightings. Our results reproduce key behaviors reported in the original DeepClean study for linear noise subtraction, while highlighting practical challenges associated with nonlinear noise removal under constrained data conditions.

I. INTRODUCTION

Gravitational-wave astronomy has rapidly evolved since the first direct detection of a binary black hole merger, opening a new observational window into the universe. Current ground-based interferometers, including LIGO, Virgo, and KAGRA, have detected dozens of compact binary coalescences involving black holes and neutron stars. However, astrophysical population studies suggest that these detections represent only a small fraction of the true event rate, with many signals remaining hidden below the detectors' noise floors.

The sensitivity of gravitational-wave detectors is fundamentally constrained by a combination of irreducible noise sources and removable environmental disturbances. At high frequencies, quantum shot noise dominates and is well understood within existing noise budgets. At lower frequencies—particularly below $\sim 100\,\mathrm{Hz}$—the measured strain often exceeds modeled noise contributions, indicating the presence of poorly understood environmental and technical couplings. This low-frequency regime is especially important for high-mass binary black hole systems, whose signals spend little time in band before merger. Consequently, improving sensitivity in this region can significantly expand the observable volume of the universe.

Traditional noise-subtraction methods rely on linear regression techniques, such as Wiener filtering, which use auxiliary "witness" sensors to estimate and subtract noise contributions from the strain channel. These methods are highly effective when noise couples linearly and remains stationary over time. In practice, however, many disturbances arise from nonlinear interactions between environmental motion and interferometer control systems, or from time-varying coupling mechanisms. Such effects limit the effectiveness of purely linear approaches and motivate the exploration of more flexible noise-subtraction strategies.

Deep learning offers a natural framework for addressing these challenges. By learning nonlinear mappings directly from data, neural networks can model complex and nonstationary relationships without requiring explicit physical descriptions of the underlying couplings. The DeepClean [1] framework applies this idea to gravitational-wave data by training a convolutional neural network to predict noise contributions using witness channels that are guaranteed to be free of astrophysical signals.

The objective of this project is to study the performance of the DeepClean approach under realistic constraints imposed by publicly available data. Specifically, we reproduce and extend aspects of the original DeepClean methodology using simulated environmental noise injected into cleaned gravitational-wave strain. We evaluate the framework on both linear and nonlinear noise scenarios, analyze the role of combined time- and frequency-domain loss functions, and assess the extent to which DeepClean can recover known noise-subtraction behavior in a controlled setting.

II. DATASET

A. Data Availability and Constraints

The original DeepClean study was developed and evaluated using raw gravitational-wave strain data and a large set of auxiliary witness channels recorded during LIGO observing runs. However, the majority of these auxiliary channels are not publicly available. Public data releases through the Gravitational Wave Open Science Center (GWOSC) provide cleaned strain data along with only a limited subset of witness channels, which restricts direct reproduction of the original DeepClean training pipeline.

This limitation particularly affects data from the O2 observing run, where publicly available strain has already undergone standard noise subtraction. As a result, realistic environmental noise couplings cannot be directly learned from the available data alone. To address this constraint, we adopt a simulation-based strategy that enables controlled evaluation of DeepClean while remaining consistent with the philosophy of the original framework.

B. Simulated Noise Injection

To construct a suitable dataset for training and evaluation, we begin with cleaned gravitational-wave strain data and algorithmically inject simulated environmental noise. This approach allows us to introduce known noise sources with controlled coupling behavior, enabling quantitative assessment of the network's performance.

The injected noise includes multiple components designed to reflect realistic disturbances observed in gravitational-wave detectors. These components include broadband noise, narrowband power mains interference centered at 60 Hz, and sideband structures that emulate modulation effects arising from control system interactions. In addition, real jitter noise measured from LIGO's pre-stabilized laser auxiliary channels is incorporated to preserve realistic temporal structure. This methodology follows the validation philosophy used in the original DeepClean study, in which simulated noise is employed to benchmark subtraction performance under known conditions.

C. Dataset Construction

The dataset is constructed from contiguous 1024 s segments of strain data. For each validation scenario, one segment is used for training the neural network, while a separate, non-overlapping segment is reserved for testing. This separation ensures that model performance is evaluated on data that was not seen during training, reducing the risk of overfitting.

We consider both O2 and O3 observing run contexts depending on the validation case. Linear noise subtraction experiments focus on O2-era data with injected jitter noise, while nonlinear noise studies are motivated by features observed during O3, such as power mains sidebands. All injected noise is added directly to the strain data prior to preprocessing.

D. Data Preprocessing

Prior to training, all data channels undergo preprocessing to ensure numerical stability and compatibility with the neural network architecture. Because the model maps input time series to output time series, all witness channels are resampled to match the sampling frequency of the gravitational-wave strain channel.

Both strain and witness channels are normalized to have zero mean and unit variance. This step prevents channels with larger absolute amplitudes from dominating the loss function and mitigates numerical issues arising from the extremely small magnitude of gravitational-wave strain, which typically lies in the range 10^{-21} to 10^{-23}. Importantly, this normalization procedure is fully invertible, allowing the network's predicted noise output to be transformed back into physical units during post-processing.

III. METHODS

The code used to implement the DeepClean architecture, simulate environmental noise, and perform training and evaluation for this project is publicly available at

`https://github.com/LemonCut/DeepClean-LIGO`

A. Model Architecture

The DeepClean framework employs a one-dimensional convolutional neural network (CNN) designed to perform time-series regression on auxiliary witness channels. The model follows a fully convolutional autoencoder-style architecture, in which successive convolutional layers compress the input representation, followed by transposed convolutional layers that reconstruct a prediction of the noise contribution in the gravitational-wave strain.

The network takes as input a set of synchronized witness channels that monitor environmental and instrumental conditions near the interferometer. These channels are explicitly chosen to contain no astrophysical gravitational-wave signal, ensuring that any learned mapping cannot distort genuine signals present in the strain channel. The network output is a time-domain estimate of the noise component coupled into the strain, which can then be subtracted from the measured data.

Convolutional layers are well suited for this task because they naturally capture local temporal correlations while allowing deeper layers to model longer-range dependencies. Nonlinear activation functions enable the network to learn both linear and nonlinear coupling mechanisms between witness channels and the strain. A schematic of the model architecture is shown in Fig. 1.

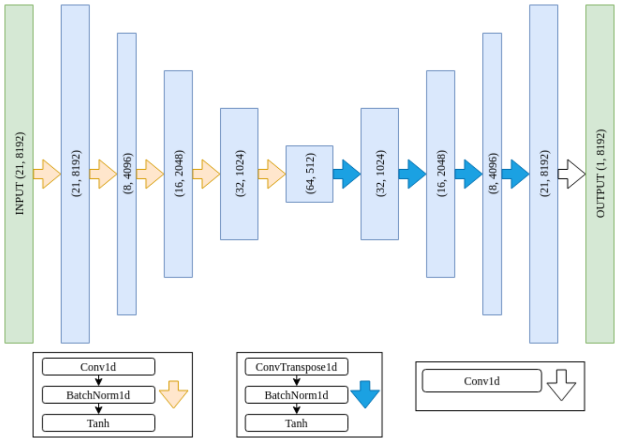

FIG. 1. Schematic of the DeepClean convolutional neural network architecture. Witness-channel time series are processed through successive convolutional and transposed convolutional layers to predict the noise contribution in the gravitational-wave strain.

B. Loss Function

The measured gravitational-wave strain can be expressed as

$$h(t) = s(t) + n(t), \tag{1}$$

where $s(t)$ represents the astrophysical gravitational-wave signal and $n(t)$ denotes environmental and instrumental noise.

The DeepClean model predicts a noise estimate $F(t)$ using only auxiliary witness channels, which do not contain astrophysical signal content. The residual strain after subtraction is therefore

$$r(t) = h(t) - F(t). \tag{2}$$

Because $F(t)$ is computed exclusively from witness data, the subtraction process cannot remove or distort the true gravitational-wave signal $s(t)$.

To quantify performance in the time domain and target spectral line noise, the DeepClean authors defined a mean squared error loss

$$\mathcal{L}_{\mathrm{MSE}} = \frac{1}{N} \sum_{i=1}^{N} r^2(t_i), \tag{3}$$

where N is the number of time samples.

To target broadband noise suppression, the authors additionally define a frequency-domain loss based on the amplitude spectral density (ASD) of the residual. This ASD loss is given by

$$\mathcal{L}_{\mathrm{ASD}} = \frac{1}{M} \sum_{j=1}^{M} \left| \frac{\mathrm{ASD}_{rr}(f_j)}{\mathrm{ASD}_{hh}(f_j)} \right|, \tag{4}$$

where $\mathrm{ASD}_r(f_j)$ is the amplitude spectral density of r(t) evaluated at frequency bin f_j. This is weighted by the reciprocal of the amplitude spectral density of the original strain to normalize each frequency bin and ensure certain bins do not dominate this loss term.

The total training loss is a weighted combination of the two components,

$$\mathcal{L} = w\,\mathcal{L}_{\mathrm{ASD}} + (1-w)\,\mathcal{L}_{\mathrm{MSE}}, \tag{5}$$

where w controls the relative emphasis on frequency-domain versus time-domain objectives.

C. Training Pipeline

Training is performed on fixed-length segments of strain and witness data using a sliding-window approach. For each training segment, witness-channel time series are passed through the network to generate a predicted noise signal. The loss is computed by comparing this prediction to the known injected noise component in the strain channel.

The model is trained exclusively on witness channels; the strain signal is used only in the computation of the loss function. This design ensures that the network does not have direct access to astrophysical signal content during inference. To prevent overfitting, training and testing are conducted on non-overlapping data segments.

Optimization is carried out using the Adam optimizer with a fixed learning rate. Training typically converges within a small number of epochs, requiring only a few minutes on standard hardware. Once trained, inference is computationally efficient, allowing noise subtraction to be performed in seconds for each data segment. A schematic of the training and inference pipeline is shown in Fig. 2.

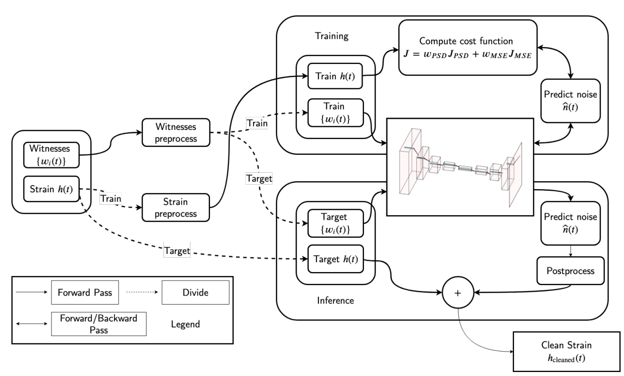

FIG. 2. Training and inference pipeline for DeepClean. Witness channels are preprocessed and passed through the neural network to predict noise, which is subtracted from the strain channel during inference.

IV. RESULTS

A. Linear Noise Validation

To evaluate DeepClean on a known linear noise source, we trained the model on 1024 s of gravitational-wave strain data from the O2 observing run with injected jitter noise from LIGO's pre-stabilized laser system. The trained network was then applied to a separate, non-overlapping 1024 s segment to assess generalization performance.

Model behavior was examined across different values of the loss weighting parameter w. In regimes where the loss was dominated by the mean squared error term (low w), the network successfully reduced time-domain noise but exhibited signs of overfitting, particularly when w approached zero. Introducing a nonzero contribution from the frequency-domain ASD loss improved stability and reduced overfitting by encouraging consistency in the spectral structure of the residual.

Figure 3 shows representative strain spectra before and after noise subtraction for different values of w. For intermediate values of w, the network effectively subtracts the injected linear noise while preserving the overall spectral shape of the strain. These results qualitatively reproduce the behavior reported in the original DeepClean study for linearly coupled noise sources.

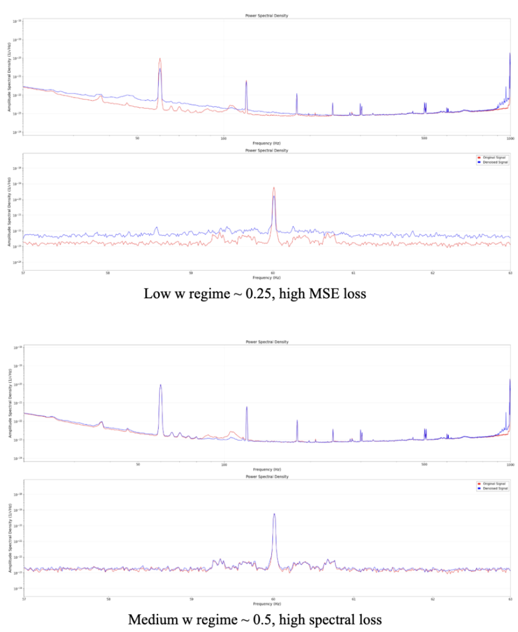

FIG. 3. Linear noise subtraction results for injected jitter noise. Spectral performance is shown for different values of the loss weighting parameter w, illustrating the tradeoff between time-domain accuracy and spectral regularization.

B. Empirical Behavior of the Loss Weighting

Further inspection of training behavior reveals that extremely small values of w lead to rapid convergence accompanied by poor generalization, indicating overfitting to time-domain features. Incorporating even a modest contribution from the ASD loss acts as a form of regularization, constraining the model to produce physically reasonable spectral behavior.

This empirical observation highlights the importance of balancing time- and frequency-domain objectives when training neural networks for noise subtraction. In practice, optimal performance is achieved in a regime where both loss components contribute meaningfully to the training objective.

C. Nonlinear Noise Validation

We next consider DeepClean's performance on nonlinearly coupled noise. In this scenario, the dominant disturbance is a 60 Hz power mains signal that couples linearly into the strain, accompanied by sidebands arising from nonlinear modulation by the alignment sensing and control system. Such sideband structures are known to challenge classical linear noise-subtraction techniques.

In principle, the nonlinear representational capacity of a neural network should enable it to learn both the primary spectral line and its associated sidebands. The original DeepClean study demonstrated successful subtraction of these nonlinear features using a rich set of auxiliary channels and real detector couplings. Figure 4 reproduces representative results from that work, illustrating substantial suppression of both the central line and the nonlinear sidebands.

In our simulation-based implementation, however, we were unable to reproduce this level of nonlinear noise subtraction consistently. While the dominant 60 Hz component was partially reduced, residual sideband structure persisted in the cleaned strain. This limitation is likely attributable to the simplified nature of the simulated nonlinear coupling and the restricted diversity of available witness channels in publicly accessible data.

These findings emphasize that, although DeepClean is theoretically well suited for nonlinear noise regression, practical performance depends sensitively on the realism of the coupling mechanisms and the availability of comprehensive auxiliary sensor information.

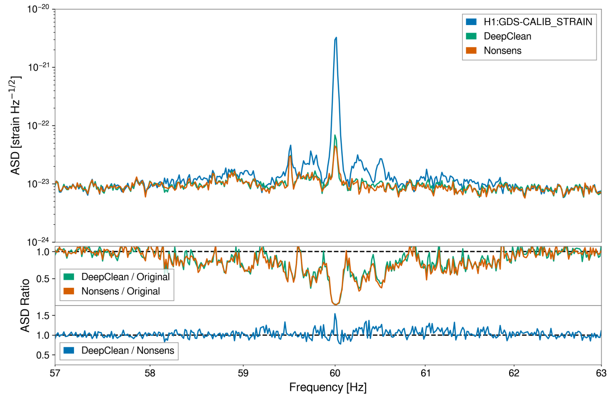

FIG. 4. (from cited paper) : Nonlinear noise subtraction results for simulated 60 Hz power mains interference with sidebands. While the dominant spectral line is reduced, residual nonlinear features persist.

D. Discussion

Overall, the results demonstrate that DeepClean reliably reproduces expected behavior for linearly coupled noise under controlled simulation conditions. The observed sensitivity to the loss weighting parameter emphasizes the role of frequency-domain regularization in stabilizing training.

For nonlinearly coupled noise, performance is more limited, reflecting both the challenges inherent in modeling nonlinear interactions and the constraints imposed by simulated data and limited witness channels. These results are consistent with the broader conclusion that access to comprehensive auxiliary sensor data is a critical factor in the success of deep-learning–based noise subtraction frameworks.

V. CONCLUSION

In this project, we investigated the DeepClean framework for noise subtraction in gravitational-wave data, focusing on its performance under realistic constraints imposed by publicly available datasets. By adopting a simulation-based strategy, we were able to evaluate the model's behavior on both linearly and nonlinearly coupled noise sources while remaining consistent with the methodology of the original DeepClean study.

Our results demonstrate that DeepClean reliably reproduces expected noise-subtraction behavior for linearly coupled disturbances. In particular, the model effectively removes injected jitter noise when trained with an appropriate balance between time-domain and frequency-domain loss terms. The inclusion of an amplitude spectral density loss was found to play a critical role in stabilizing training and mitigating overfitting, acting as a form of spectral regularization.

For nonlinearly coupled noise, performance was more limited. While the dominant power mains line was partially suppressed, residual sideband structure persisted. These results highlight the challenges associated with learning nonlinear coupling mechanisms in the absence of rich auxiliary sensor data and realistic coupling dynamics. They further emphasize that access to comprehensive witness channels is likely essential for achieving the full potential of deep-learning–based noise subtraction in operational detector settings.

Overall, this study supports the conclusion that DeepClean is a promising and flexible framework for gravitational-wave noise mitigation, particularly for linear and mildly nonlinear noise sources. Future work could extend this analysis by incorporating more realistic nonlinear simulations, leveraging a broader set of auxiliary channels, or applying the framework to low-latency data streams in upcoming observing runs.

ACKNOWLEDGMENTS

We thank Prof. J Duarte and TA D. Primosch for guidance throughout this project. We also acknowledge the LIGO Scientific Collaboration and the Gravitational Wave Open Science Center for providing access to public gravitational-wave data. This work was inspired by the DeepClean framework introduced by Ormiston et al., and we thank the authors for making their methodology publicly available.

[1] R. Ormiston, T. Nguyen, M. Coughlin, R. X. Adhikari, and E. Katsavounidis, "Noise reduction in gravitational-wave data via deep learning," *Phys. Rev. Research* **2**, 033066 (2020).

[2] Gravitational Wave Open Science Center, `https://www.gw-openscience.org`

[3] J. Aasi *et al.* (LIGO Scientific Collaboration), "Advanced LIGO," *Classical and Quantum Gravity* **32**, 074001 (2015).